# A Lévy-driven rainfall model with applications to futures pricing

Ragnhild C. Noven  Almut E. D. Veraart  Axel Gandy

Department of Mathematics, Imperial College London

arXiv:1403.7406v1 [stat.ME] 28 Mar 2014

**Abstract**

We propose a parsimonious stochastic model for characterising the distributional and temporal properties of rainfall. The model is based on an integrated Ornstein-Uhlenbeck process driven by the Hougaard Lévy process. We derive properties of this process and propose an extended model which generalises the Ornstein-Uhlenbeck process to the class of continuous-time ARMA (CARMA) processes. The model is illustrated by fitting it to empirical rainfall data on both daily and hourly time scales. It is shown that the model is sufficiently flexible to capture important features of the rainfall process across locations and time scales. Finally we study an application to the pricing of rainfall derivatives which introduces the market price of risk via the Esscher transform. We first give a result specifying the risk-neutral expectation of a general moving average process. Then we illustrate the pricing method by calculating futures prices based on empirical daily rainfall data, where the rainfall process is specified by our model.

## 1. Introduction

A typical rainfall time series has several properties which are difficult to capture in a simple statistical model, including a heavily skewed marginal distribution which is distinctly non-Gaussian, a large proportion of zero values, and frequent large fluctuations. Thus there is a need for specialised models for rainfall which can capture the unique characteristics of this type of process. The existing literature on modelling rainfall is large and spread over fields such as hydrology, atmospheric sciences, environmental risk analysis and statistics. Onof et al. (2000) classifies the different approaches into four categories: meteorological models based on large sets of differential equations, multi-scale models concerned with the spatial evolution of rainfall, statistical models that capture spatial and temporal trends, and finally stochastic process models which make simple assumptions in order to remain parsimonious. In the following we will focus on models of the last category.

Many of the early attempts at modelling rainfall use a simple model which represents the rainfall occurrence process as a two-state Markov chain, and models the intensity of rainfall with a Gamma distribution (Katz, 1977; Chin, 1977; Woolhiser and Roldán, 1982; Coe and Stern, 1982). This model is easy to interpret and enables the direct use of likelihood methods for fitting. However, it makes several restrictive assumptions on the rainfall process, and may require a high-order Markov chain with many parameters to capture observed temporal dependence.

There is also a large literature on modelling rainfall for hydrological applications based on a form of the Poisson-cluster model, first developed by Rodriguez-Iturbe et al. (1987) and Cox and Isham (1988). This model is based on a hierarchical structure, with a primary

Poisson Process controlling the arrival of storms and a secondary process generating cells from each storm, which then deposit rainfall. There have been numerous extensions of the Poisson-cluster model aimed at particular applications (e.g. Cowpertwait et al., 2007; Northrop, 1998). As remarked by Onof et al. (2000), developments of the cluster model face a challenging trade-off between the inclusion of more features and their mathematical tractability.

The contribution of the present paper is twofold. First, we develop a parsimonious model that captures the distributional and time-dependent features of the observed rainfall time series. We show how the model relates to the framework of Lévy-driven, continuous-time ARMA (CARMA) processes, and based on this connection we develop a suitable fitting method which is illustrated using empirical rainfall data. Second, we derive a formula specifying the risk-neutral distribution of a general class of Lévy-driven stochastic processes, which includes our model as a special case. We then use this result to calculate prices for rainfall futures based on our model.

Rainfall derivatives were introduced at the Chigago Merchantile Exchange (CME) in 2010, as a recent addition to the class of weather-related products. These products have a large potential market in all economic sectors which depend on favourable weather conditions, such as farming and energy development. The literature on rainfall derivatives pricing is currently limited. The papers by Leobacher and Ngare (2011) and Carmona and Diko (2005) both use the utility indifference approach to pricing, based on, respectively, the Markov-Gamma and the Poisson-cluster model. López Cabrera et al. (2013) use a version of the daily rainfall model by Wilks (1998). They fit simulated monthly rainfall totals to a normal inverse Gaussian distribution which is then transformed according to a risk premium via the Esscher transform (Esscher, 1932). Benth and Šaltytė Benth (2013) also use the Esscher transform for pricing, but base their underlying rainfall model on an independent increment process.

Our proposed model has the advantage of not making assumptions about temporal independence or Markovianity of the rainfall process increments. This results in a model with a flexible autocorrelation structure, which is particularly relevant for data on finer time scales. This flexibility is illustrated by fitting the model to hourly rainfall data. We also compare our model to that given in Wilks (1998), which was used in López Cabrera et al. (2013) for pricing rainfall futures based on daily data from Detroit.

This paper is structured as follows: Section 2 discusses characteristic features of the rainfall process in light of data from different locations and time scales. Section 3 presents the rainfall model and shows how it fits into the continuous-time ARMA (CARMA) model framework. Section 4 gives details on the fitting method, and Section 5 investigates the model performance using empirical data. In Section 6 we derive a method for pricing rainfall derivatives based on our model.

## 2. Characteristics of observed rainfall

In this section we motivate the structure of our model by illustrating some of the characterising features of rainfall time series. We base this illustration on two data sets which will be used throughout this paper: the first consists of hourly accumulated rainfall amounts at Heathrow (UK) over the years 2010-2012, provided by the UK Meteorological Office (2012). The second data set gives daily accumulated rainfall amounts in Detroit (US) over the years 1980-2010, provided by Bloomberg Professional Service.

Figure 1 shows the rainfall time series for both locations over the last three years of observations. These plots illustrate that the rainfall process is subject to sudden transitions between

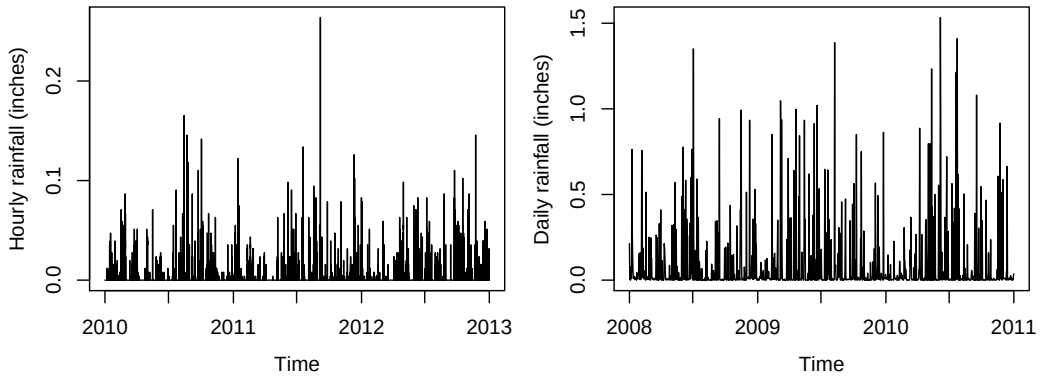


Figure 1: Time series of empirical rainfall over three years for Heathrow (left) and Detroit (right).

| Location | Heathrow | Heathrow | Detroit |
|---|---|---|---|
| Time scale | Hourly | Daily | Daily |
| Zero values (%) | 91.85 | 57.53 | 48.12 |

Table 1: Percentage of rainfall measurements equal to zero.

periods with little or no rain and periods of higher intensity, causing the large spikes in the graph.

There is also a large amount of zero values in both data sets, as shown in Table 1. Because the data is rounded to the nearest unit of measurement (0.1mm for the Heathrow data and 3/100 inch for the Detroit data), these zero values represent all data points with a value smaller than this unit. The proportion of zero values is dependent on the measurement time scale, with fewer zero values for the daily data. This is because periods with no rainfall must persist throughout the measurement time interval in order to induce a value of zero in the data.

Figure 2 shows histograms of the non-zero (i.e. positive-valued) data, which illustrates the non-normality and skewness of the empirical distributions. In the hourly data the skewness is more pronounced, giving a coefficient of 10.7, compared with 4.5 for the daily data. This fits with the general observation that measuring accumulated rainfall on larger time scales has a smoothing effect, which makes characterising features such as large skewness and frequent zero values less evident.

The empirical autocorrelation functions of the Heathrow and Detroit rainfall time series are shown in Figure 3. For the hourly Heathrow data there is clearly a non-trivial autocorrelation part which decays smoothly up to lag 10. For the daily Detroit data the autocorrelation function (ACF) decays steeply after lag 1, indicating that there is less relevant time-dependence in this rainfall process, as one would expect from the daily time scale. As will be seen in Section 4, these differences in the autocorrelation structure lead us to fit models of different orders to the two data sets.

## 3. Rainfall Model

In this section we present the basic structure of our model and relate its properties to the observed rainfall dynamics. Based on this we construct a model extension where the rainfall

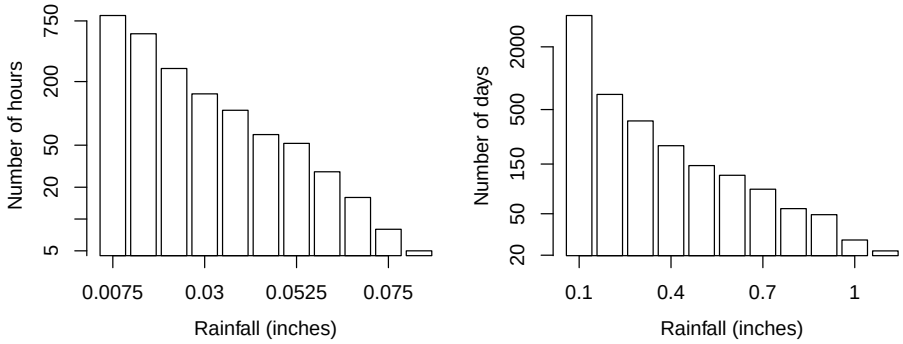


Figure 2: Histogram of non-zero rainfall for Heathrow (left) and Detroit (right), with counts on a log scale.

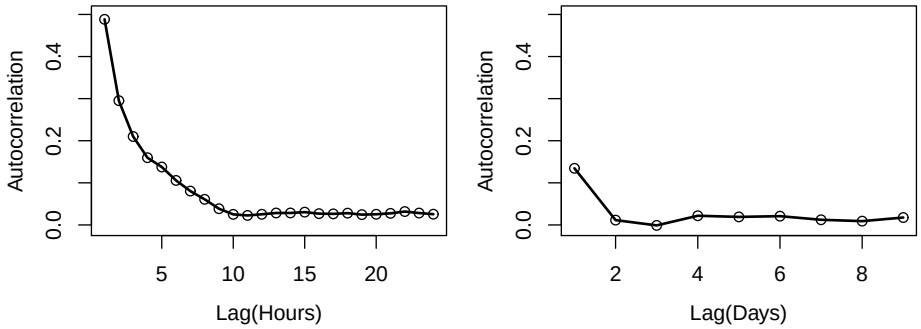


Figure 3: Empirical rainfall autocorrelation functions for Heathrow (left) and Detroit (right).

intensity process belongs to the class of continuous-time ARMA (CARMA) models.

### 3.1. Primary model structure

We propose to model the accumulated rainfall $R$ by

$$R(t_i) - R(t_{i-1}) = S(t_{i-1})(Y(t_i) - Y(t_{i-1})), \tag{1}$$

where $0 = t_0 < t_1 < \ldots < t_n$ are discrete measurement times such that $t_i - t_{i-1} = \delta$, and $S$ is a deterministic seasonal component. We define the continuous-time stochastic process $(Y(t))_{t\geq 0}$ as the integral of a Lévy-driven Ornstein-Uhlenbeck (OU) process $(X(t))_{t\geq 0}$ (Barndorff-Nielsen and Shephard, 2001), i.e.

$$\begin{aligned} Y(t) &= \int_0^t X(s)\, ds, \\ X(s) &= X(0)e^{-\lambda s} + \int_0^s e^{-\lambda(s-v)}\, dL(v), \end{aligned} \tag{2}$$

where $\lambda$ is a positive parameter and $(L(t))_{t\geq 0}$ is a Lévy subordinator. We interpret $X(s)$ as the instantaneous rainfall intensity at time $s$, and so $Y(t)$ measures the accumulated rainfall over the time interval $[0, t]$ - up to the seasonal adjustment given by $S$. We let $X(0)$ be a random variable which is independent of $(L(t))_{t\geq 0}$ and has distribution

$$X(0) \stackrel{d}{=} \int_0^\infty e^{-\lambda v}\, dL(v),$$

making the resulting OU process $X$ strictly stationary.

Using the Fubini theorem for stochastic integrals, we can exchange the order of integration in the definition of $Y$ and obtain a simpler expression in terms of a single stochastic integral:

$$Y(t) = X(0)\left(\frac{1-e^{-\lambda t}}{\lambda}\right) + \int_0^t \frac{1-e^{-\lambda(t-v)}}{\lambda}\, dL(v). \tag{3}$$

For our particular rainfall model we let $(L(t))_{t\geq 0}$ be a compound Poisson process with Gamma-distributed jumps, called the Hougaard process (Lee and Whitmore, 1993; Grigelionis, 2011). This means that $L$ is a pure-jump Lévy process, more specifically a subordinator. The choice of a pure-jump Lévy process is motivated by the intermittent behaviour of the observed rainfall process, in particular the abrupt switches from exact zero to large positive values, which are modelled by jumps in the driving process $L$.

The marginal distribution of $L(1)$ is a member of the Tweedie distribution family (Jørgensen, 1997), which was used by Dunn (2004) to model the monthly rainfall in Australia. The Tweedie distribution is attractive for use in rainfall modelling because it has an atom at zero, which can be interpreted as the probability of observing no rainfall during a certain time period. In the following we will parameterise $L(1)$ as a Tweedie random variable, which has parameters $(\mu, \phi, p)$ such that

$$\begin{aligned} \mathrm{E}(L(1)) &= \mu, \\ \mathrm{Var}(L(1)) &= \phi\mu^p. \end{aligned} \tag{4}$$

The stochastic process $X$ defined in (2), which represents the rainfall intensity, has an interpretation in terms of the physical dynamics of the rainfall process. In this interpretation the jumps of the driving process $L$ represent the arrival of storm events, generating a jump in the intensity of random size. As the storm dissipates this intensity decays smoothly towards zero at a rate determined by the parameter $\lambda$ in the OU process. Considering this interpretation we see some similarity to the Poisson-cluster models mentioned in the introduction, which are based on the idea of storms arriving according to a Poisson process. The mathematical setup is rather different, however, as Poisson-cluster models have a hierarchical structure, as opposed to our stochastic integral specification in (3).

From (1) it is clear that the discrete-time process $(\triangle Y(t_i))_{i=1,\ldots,n}$ given by

$$\triangle Y(t_i) := Y(t_i) - Y(t_{i-1})$$

should have features resembling those of the deseasonalised empirical rainfall. As shown in Section 5, Figures 4 and 5, the empirical marginal distribution of our rainfall data is well approximated by the marginal distribution of $\triangle Y$. In fact, this was the motivation for specifying the driving process $L$ to be the Hougaard process.

When it comes to approximating the empirical autocorrelation structure, the suitability of the model is not so clear. As the following proposition shows, the autocovariance function $C_{\triangle Y}(h)$ is restricted to take the form of an exponential decay.

**Proposition 1.** *The autocovariance function of $\triangle Y$ is given by*

$$\begin{aligned} C_{\triangle Y}(0) &= \frac{\phi\mu^p}{\lambda^3}(e^{-\lambda\delta} + \lambda\delta - 1), \\ C_{\triangle Y}(h) &= \frac{\phi\mu^p}{2\lambda^3}(e^{-\lambda\delta} - 1)^2 e^{-\lambda(h-1)\delta}, \quad h \geq 1. \end{aligned} \tag{5}$$

The proof of Proposition 1 is given in the Appendix. As illustrated in Figure 3, the empirical autocovariance functions do not necessarily take this simple form. This restrictive form of the autocovariance function motivates the following extension of the model.

## 3.2. Extension to CARMA process

In this subsection we consider an extension of our model which admits a more flexible autocovariance structure. This extension is based on generalising the Ornstein-Uhlenbeck process $X$ in (2) to a continuous-time ARMA (CARMA) process. We first give a brief overview of the construction of Lévy-driven CARMA processes, and then show how the extension of $X$ is obtained.

### 3.2.1. CARMA processes

A CARMA processes is a continuous-time analogue of the discrete-time ARMA process. Here we will consider Lévy-driven CARMA processes (Brockwell, 2001; Brockwell and Lindner, 2009). To illustrate the correspondence to the discrete-time setting, we start by considering the ARMA$(p, q)$ process $(V_n)$ defined by the difference equation

$$a(B)V = b(B)L,$$

where $B$ is the backward shift operator, $L$ is a white noise sequence and $a, b$ are polynomials given by

$$\begin{aligned} a(x) &= x^p + a_1 x^{p-1} + \ldots + a_p, \\ b(x) &= b_0 + b_1 x + \ldots + b_q z^q. \end{aligned}$$

We can consider formally replacing $B$ with the differential operator $D$ to obtain a stochastic differential equation (SDE) for the CARMA(p,q) process $V$ driven by the process $L$.

This SDE will contain expressions of the form $D^k L$, which may not be well-defined. Therefore it is customary to consider an equivalent definition of CARMA processes via the state-space representation. This representation defines the observation and state equations

$$\begin{aligned} V(t) &= \mathbf{b}^T \mathbf{Z}(t), \\ d\mathbf{Z}(t) - A\mathbf{Z}(t)dt &= \mathbf{e}\, dL(t), \end{aligned}$$

where $\mathbf{b}$ is the vector of coefficients of $b(x)$, $A$ is the matrix

$$\begin{pmatrix} 0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ -a_p & -a_{p-1} & -a_{p-2} & \ldots & -a_1 \end{pmatrix},$$

where $a_i$ is the ith coefficient of $a(x)$, $\mathbf{e}$ is the pth unit vector and $L$ is the driving Lévy process. Provided all eigenvalues of $A$ have negative real parts, the SDE for $Z$ can be solved to give the following expression for $(V(t))_{t\geq 0}$:

$$V(t) = \mathbf{b}^T e^{At}\mathbf{e} V(0) + \int_0^t \mathbf{b}^T e^{A(t-u)}\mathbf{e}\, dL(u), \tag{6}$$

where $V$ is strictly stationary. We can now obtain a representation of the CARMA process $V$ that extends the Lévy-driven OU process (Brockwell, 2004). Let $\{\alpha_i\}$ be the eigenvalues of the matrix $A$, or equivalently the roots of the polynomial $a(x)$, with corresponding eigenvectors

$$(1, \alpha_i, \ldots, \alpha_i^{p-1}).$$

Then we can obtain the spectral expansion

$$\mathbf{b}^T e^{A(t-u)}\mathbf{e} = \sum_{i=1}^{p} \frac{b(\alpha_i)}{a'(\alpha_i)} e^{\alpha_i(t-u)}.$$

Substituting this expansion into (6) gives

$$V(t) = V(0)\left(\sum_{i=1}^{p} \frac{b(\alpha_i)}{a'(\alpha_i)} e^{\alpha_i t}\right) + \int_0^t \left(\sum_{i=1}^{p} \frac{b(\alpha_i)}{a'(\alpha_i)} e^{\alpha_i(t-u)}\right) dL(u),\ t \geq 0. \tag{7}$$

If we now set $p = 1$ and let $a(z) = z - \lambda, b(z) = 1$, we recover the OU process $X$ given in (2), with $\lambda = -\alpha_1$. Thus CARMA processes generalise Ornstein-Uhlenbeck processes, and this motivates the extended model described in the following.

### 3.2.2. Extended model

We define the extended model of order $p$ by

$$R(t_i) - R(t_{i-1}) = S(t_{i-1})(Y(t_i) - Y(t_{i-1})),$$

where

$$Y(t) = \int_0^t X(s)\, ds, \tag{8}$$

$$X(s) = \left(\sum_{i=1}^{p} w_i X_i(0) e^{-\lambda_i s}\right) + \int_0^s \left(\sum_{i=1}^{p} w_i e^{-\lambda_i(s-v)}\right) dL(v) := \sum_{i=1}^{p} w_i X_i(s), \tag{9}$$

with $\sum_i w_i = 1$. Here $X$ is a CARMA process of the form given in (7), where the polynomial $b$ is of order $p-1$, and its coefficients can be found by solving $b(-\lambda_i)/a'(-\lambda_i)$ for $w_i$, with $b_{p-1} = 1$. Similarly to the OU case, each $X_i(0)$ is chosen to be independent of $(L_t)_{t\geq 0}$, with

$$\sum_{i=1}^{p} w_i X_i(0) e^{-\lambda_i s} \overset{d}{=} \int_0^\infty \left(\sum_{i=1}^{p} w_i e^{-\lambda_i(s+v)}\right) dL(v).$$

This extension of $X$ can also be seen as a mixture of dependent OU processes $X_i$, driven by the same subordinator $L$.

As we will see in equation (10), the autocovariance function $C_{\triangle Y}$ with $Y$ defined in (8) is a mixture of exponential decays with separate rates $\lambda_i$. Hence we can get more complex autocovariance structures by increasing the order $p$ of the CARMA process $X$.

## 4. Fitting procedure

In this section we discuss a fitting approach for the extended model defined in Section 3.2.2. The fitting is done in two parts, first we estimate the autocovariance parameters via the CARMA representation, and then we find moment-based estimates of the driving Lévy process parameters.

In the following discussion we assume that the deterministic seasonality component $S$ has already been fitted using a standard technique; details of our approach are given at the end of the section. We then get from (1) that

$$\frac{R(t_i) - R(t_{i-1})}{S(t_{i-1})} = Y(t_i) - Y(t_{i-1}),$$

so we can fit $\triangle Y$ to the discrete observations $\triangle R/S$. Thus in the following we will only consider fitting the model given by $\triangle Y$.

### 4.1. Autocovariance structure

We now show how to use the CARMA representation of the process $X$ to develop a fitting method for the parameters $\{\lambda_i, w_i\}$. This approach relies on Theorem 2 in Brockwell and Lindner (2013), which states that under certain conditions[1] on the polynomials $a$ and $b$, we have that for a causal and invertible CARMA$(p, p-1)$ process $V$, the discrete process $I_n^\Delta$ given by

$$I_n^\Delta = \int_{(n-1)\Delta}^{n\Delta} V(s)ds,$$

is a weak ARMA process. This implied ARMA process takes the form

$$\phi(B)I_n^\Delta = \theta(B)\epsilon,$$

where $\{\epsilon_i\}$ is a weak white noise sequence, i.e. the terms are uncorrelated but possibly dependent. Here the parameters of the process $L$ driving the CARMA process $V$ only affect the sequence $\{\epsilon_i\}$, not the polynomials $\phi$ and $\theta$. Furthermore, the theorem also states that there is a one-to-one correspondence between the coefficients $\{w_i, \lambda_i\}$ of $V$ and the coefficients $(\phi_i, \theta_i)$ of the corresponding ARMA$(p, p)$ process.

By using the CARMA$(p, p-1)$ representation of $X$ as defined in (9), we can write

$$\triangle Y(t_i) = \int_{t_{i-1}}^{t_i} X(s)\, ds = I_{t_i/\delta}^\delta.$$

Hence the observed increments of $Y$ can be seen as observations from the implied weak ARMA process.

We can also obtain the autocovariance of $\triangle Y$ from the integrated CARMA representation (Brockwell and Lindner, 2013, Corollary 2):

$$\begin{aligned} C_{\triangle_Y}(0) &= \sum_{\lambda_i} 2\beta(\lambda_i)\lambda_i^{-2}(e^{-\lambda\delta} - 1 + \lambda_i\delta), \\ C_{\triangle_Y}(h) &= \sum_{\lambda_i} \beta(\lambda_i)\lambda_i^{-2}(e^{-\lambda_i\delta} - 1)^2 e^{-\lambda_i(h-1)\delta}, \quad h \geq 1, \\ \beta(\lambda_i) &= \sigma^2 \frac{b(-\lambda_i)b(\lambda_i)}{a'(-\lambda_i)a(\lambda_i)}, \end{aligned} \tag{10}$$

where $\sigma^2$ is the variance of the driving process $L$ and $a, b$ are the polynomials in the CARMA representation of $X$.

We now follow Brockwell and Lindner (2013) in estimating $\{w_i, \lambda_i\}$ by minimising the weighted sum of the one-step prediction errors of the implied ARMA$(p, p)$ process, which is equivalent to minimising with respect to $\{w_i, \lambda_i\}$ due to the one-to-one correspondence. Initial

[1]The conditions are as follows: $a$ and $b$ have no common zeroes, the roots of $a$ have multiplicity 1, and $\mathrm{Im}(\lambda_i) \in (-\frac{\pi}{\delta}, \frac{\pi}{\delta})$.

values for the parameters in the minimisation can be obtained by setting the values of the autocovariance function of $\triangle Y$ for the first few lags equal to the corresponding empirical values.The estimation procedure based on minimising the prediction errors is shown to be strongly consistent by Brockwell and Lindner (2013).

We will use a CARMA$(1,0)$ model for the intensity process $X$ corresponding to the Detroit rainfall data, and a CARMA$(2,1)$ model for $X$ corresponding to the Heathrow data. These orders are chosen to be as low as possible while ensuring that the ACF of the fitted model can adequately replicate the shape of the empirical ACF. There are $2p-1$ parameters determining the ACF of the resulting model, and so it can be difficult to uniquely identify all the parameters via the minimisation procedure when $p$ is too large. This was found to be the case for the Detroit data with $p=2$. In general we expect that higher order models are more suitable for high-frequency data, which has more significant dependence structure.

For the CARMA$(1,0)$ model representing the Detroit data we have the injective mapping $\lambda \to (\phi, \theta)$ given by

$$\begin{aligned} \phi(\lambda) &= e^{-\lambda}, \\ \theta(\lambda) &= -r - \sqrt{r^2-1}, \\ r &= \frac{1-\lambda-e^{-2\lambda}(1+\lambda)}{1-2\lambda e^{-\lambda}-e^{-2\lambda}}. \end{aligned}$$

For the CARMA$(2,1)$ model, the mapping between $(w_1, \lambda_1, \lambda_2)$ and $(\phi_1, \phi_2, \theta_1, \theta_2)$ is found by numerically solving for the autocovariance function of the implied ARMA process.

### 4.2. Driving Lévy process

Having estimated $\{w_i, \lambda_i\}$ it remains to estimate the parameters $(\mu, \phi, p)$ of the driving process $L$. The parameter estimation is done by the method of moments applied to the process $\triangle Y$, which takes the form

$$\begin{aligned} \triangle Y(t) = \sum_i X_i(0)\left(\frac{e^{-\lambda_i t_{i-1}} - e^{-\lambda_i t_i}}{\lambda_i}\right) + \int_0^{t_{i-1}} \left(\sum_i \frac{e^{-\lambda_i(t_{i-1}-s)} - e^{-\lambda_i(t_i-s)}}{\lambda_i}\right) dL(s) \\ + \int_{t_{i-1}}^{t_i} \left(\sum_i \frac{1-e^{-\lambda_i(t_i-s)}}{\lambda_i}\right) dL(s), \end{aligned} \tag{11}$$

with all three terms independent due to the independent increments of $L$. It is easily shown that

$$\mathrm{E}(\triangle Y) = \sum_i \frac{\mu\delta}{\lambda_i},$$

and the variance of $\triangle Y$ is given in (10). The third moment can be calculated numerically by obtaining an expression for the characteristic function of $\triangle Y$, as shown in the Appendix.

We now replace the autocovariance parameters $\{w_i, \lambda_i\}$ in the expressions for the moments with their least-squares estimates. Comparing the theoretical moments to those of the observed increments $\triangle R/S$ gives three equations with unknowns $(\mu, \phi, p)$, which can be solved to obtain estimates for these parameters.

To obtain confidence intervals for the estimated parameters we use the block bootstrap method (Politis and Romano, 1994; Kunsch, 1989) to resample from the empirical distribution under the assumption of dependent data. In this resampling the block size has a

| Parameter | $w_1$ | $\lambda_1$ | $\lambda_2$ |
|---|---|---|---|
| Estimated value | 0.87 | 3.66 | 0.28 |
| 95% CI | $(0.83, 0.93)$ | $(2.70, 6.18)$ | $(0.21, 0.41)$ |

| Parameter | $\mu$ | $\phi$ | $p$ |
|---|---|---|---|
| Estimated value | 0.28 | 41.03 | 1.13 |
| 95% CI | $(0.22, 0.42)$ | $(33.57, 80.19)$ | $(1.06, 1.49)$ |

Table 2: Estimated parameters and confidence intervals (CI) for Heathrow rainfall data.

| Parameter | $\lambda$ | $\mu$ | $\phi$ | $p$ |
|---|---|---|---|---|
| Estimated value | 4.59 | 1.81 | 10.44 | 1.60 |
| 95% CI | $(4.18, 5.48)$ | $(1.62, 2.15)$ | $(10.06, 11.69)$ | $(1.44, 1.74)$ |

Table 3: Estimated parameters and confidence intervals (CI) for Detroit rainfall data.

geometric distribution with a specified mean value, which ensures stationarity of the resampled sequence.

In order to choose the mean block size we first generated a sample of model simulations with the parameters held fixed at their estimated values. Then we applied the block bootstrap with fixed mean block size to the simulated data sets to obtain a bootstrap confidence interval for each simulation. We repeated this procedure for several choices of the mean block size, and chose the one that maximised the coverage rate of the confidence intervals.

After obtaining a bootstrap sample as described above, we calculated 95% confidence intervals by taking the lower and upper bounds to equal, respectively, the 2.5 and 97.5 percentiles of the bootstrap sample. Tables 2 and 3 show the estimated parameters and confidence intervals for the Heathrow and Detroit data.

**Estimating the seasonality function S:** We estimated $S$ by first calculating the average monthly standard deviations over the available years. Then we fitted a Fourier Series with an annual period to these values using a least squares approach, resulting in a smooth estimate of the seasonality. For our data we found that using a Fourier series of order three resulted in a good fit.

# 5. Assessing model performance

## 5.1. Simulation study

We can simulate exactly from the process $Y$ by using the compound Poisson process representation of the driving process $L$, which gives an expression for $Y$ as a weighted sum of the jumps of $L$. Multiplying by the seasonality function $S$ then gives a simulation from the full model for the accumulated rainfall increments $\triangle R$.

Figures 4 and 5 show histograms of the simulated and empirical time series side-by-side for comparison, where the simulated time series uses the parameter estimates given in Tables 2 and 3, and has the same length as the empirical time series. For the daily Detroit data we also include the histogram of a simulation from the model given in Wilks (1998), which was fitted to the empirical data. This model was used in López Cabrera et al. (2013) as the basis for a pricing method, and so in view of the application in Section 6 it is a natural choice for comparison.

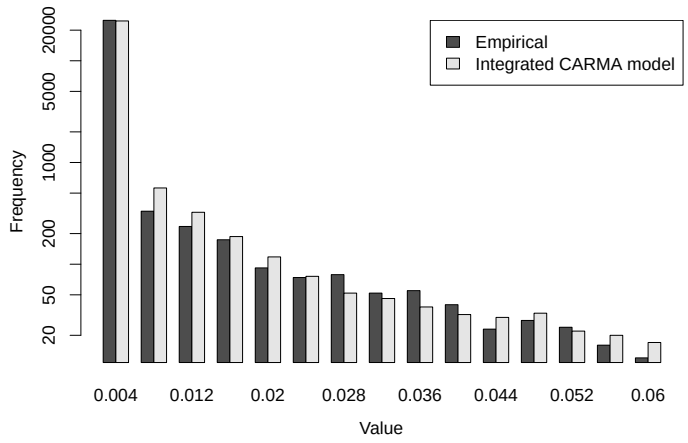


Figure 4: Histograms of empirical hourly rainfall and a simulation from the integrated CARMA model, with frequencies on a log scale.

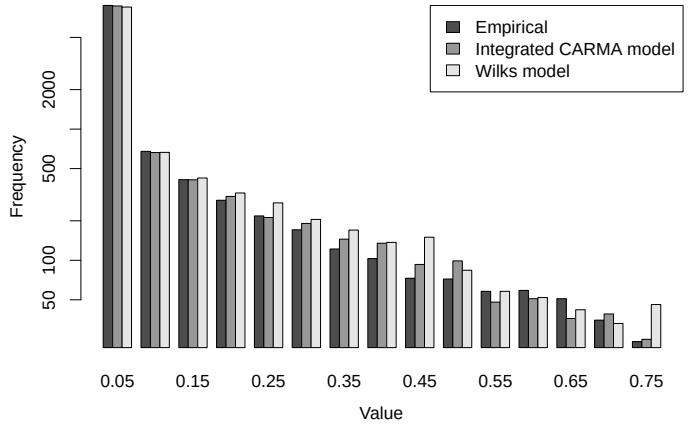


Figure 5: Histograms of empirical daily rainfall together with simulations from the integrated CARMA model and the Wilks model, with frequencies on a log scale.

We see that on both time scales the model manages to capture the characteristic shape of the rainfall distribution quite well. For the daily Detroit data the fit is somewhat better than that of the model given in Wilks (1998), especially in the lower part of the range where the majority of the data is found.

Figures 6 and 7 show path plots of the empirical data together with one particular simulation from the fitted model. For the daily Detroit data the paths look very similar, the model clearly reproduces the intermittent jump behaviour of the observed rainfall process. For the hourly Heathrow data the simulation also has intermittent jumps and periods of no activity, similar to the empirical process. However, there appears to be some clustering of time intervals with higher intensity in the empirical process which is not reflected by the model.

Figure 8 shows the autocorrelation function of the deseasonalised data along with the theoretical ACF of the fitted model for both time scales. This figure shows that our model manages to capture the autocorrelation structure of the rainfall process accurately for both hourly and daily time scales.

Table 4 shows the proportion of zero values in the empirical data, together with the corre-

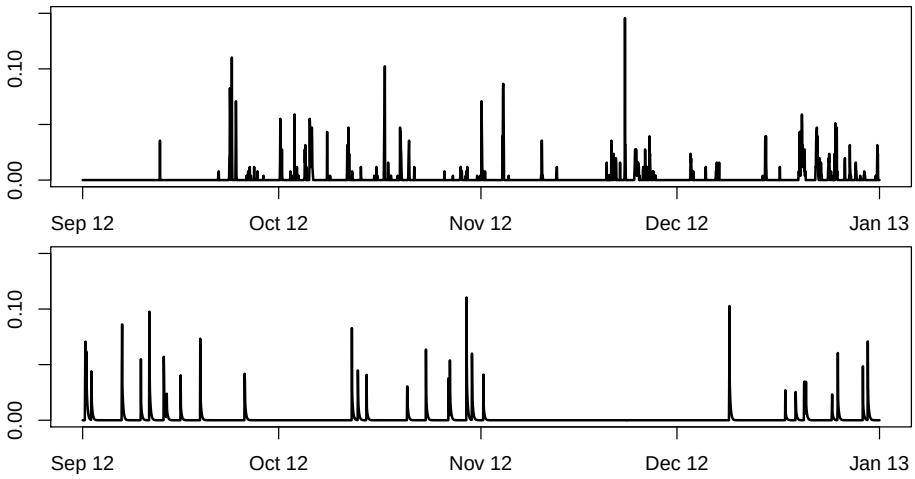


Figure 6: Path plots of hourly Heathrow empirical data (top) and simulated data (bottom).

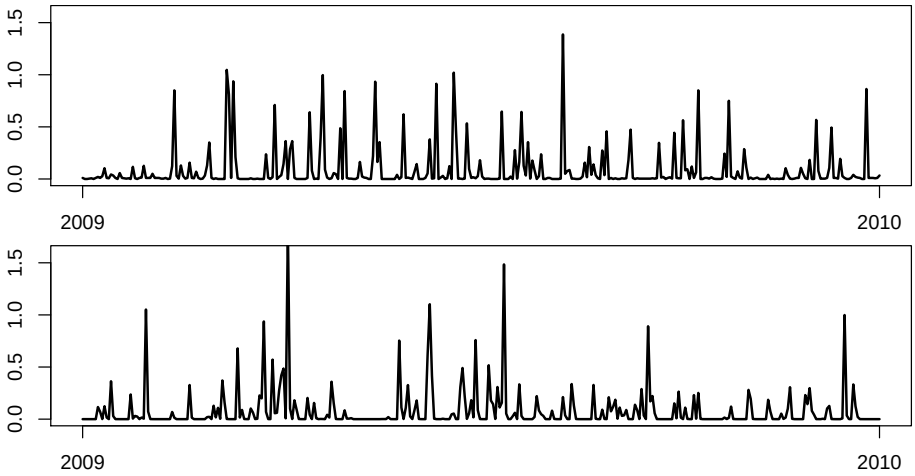


Figure 7: Path plots of daily Detroit empirical data (top) and simulated data (bottom).

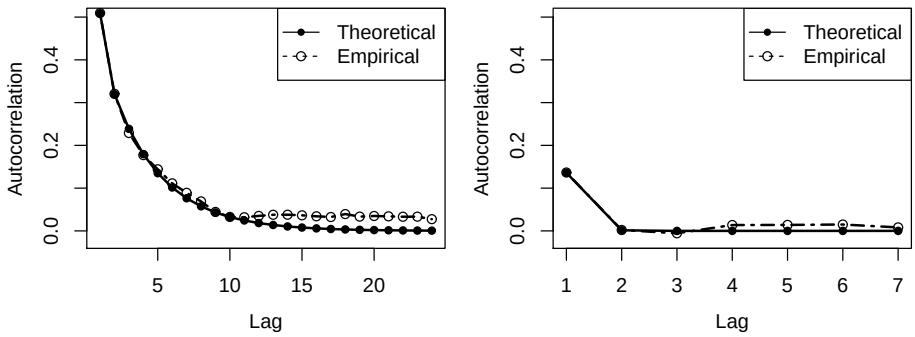


Figure 8: Theoretical and empirical ACF for Heathrow (left) and Detroit (right) rainfall.

| Location | Time scale | Type | Implied zero proportion (%) |
|---|---|---|---|
| Heathrow | Hourly | Simulation average | 93.86 |
| | | Data | 91.85 |
| Detroit | Daily | Simulation average | 60.09 |
| | | Data | 48.12 |

Table 4: Proportion of implied zero values in empirical data and corresponding simulation average.

| Model | Implied ARMA model | Random walk model | Zero model |
|---|---|---|---|
| MSE | 0.710 | 0.930 | 1.048 |

Table 5: Comparison of mean squared forecast errors (MSE).

sponding zero proportion averaged over 20 simulations. As mentioned in Section 2, the zero values in the empirical data come from data points with a value below the measurement threshold, and we performed the same rounding for the simulated data to get the *implied* zero proportion shown. We see that the simulated time series for the hourly rainfall have very similar zero proportions to the empirical data, and the daily rainfall simulations are somewhat less accurate, they overestimate the zero proportion by about 25%. In situations where it is important to match the zero proportion of the observed rainfall on all time scales exactly, a different approach for estimating the parameters of $L$ may be more appropriate. One alternative would be to use a simulation-based generalised method of moments, with one of the moment conditions specifying that the proportion of implied zero values in empirical and simulated data match.

### 5.2. Forecasting

The implied ARMA model representation which was discussed in Subsection 4.1 can be used to make short-term predictions of the rainfall. At the daily scale there does not appear to be much forecasting potential due to the rapidly decaying ACF, so we only considered forecasting the hourly Heathrow data. We remark that we do not aim to compare the forecasts from the ARMA model to those from models based on significantly more complex approaches; in particular those which are not in the stochastic process model category in the classification of Onof et al. (2000).

To evaluate the forecasting ability of our model we first performed parameter estimation based on the first two years of the hourly Heathrow data. Then we calculated 1-step-ahead forecasts for the hourly rainfall amounts in the third year by taking the linear predictors from the implied ARMA model and multiplying with the fitted seasonality function $S$. To evaluate the accuracy of forecasts from our model we compared them to the forecasts from two naive models, the first being a random walk and the second taking a constant value of zero. Table 5 shows the mean squared errors of the forecasts from all three models, where the data has been scaled to have standard deviation equal to 1.

We performed Diebold-Mariano tests (Diebold and Mariano, 1995) to compare the forecast accuracies of the implied ARMA model and the naive models, with the alternative hypothesis that our model has a better accuracy under squared error loss. This test gave a p-value less than $10^{-4}$ in both cases, thus we see that the implied ARMA model gives a significant improvement in forecasting accuracy over the random walk and zero models.

# 6. Rainfall derivative pricing

In this section we calculate prices for rainfall futures contracts based on the daily rainfall model presented in this paper. This is done by first deriving the characteristic function of a general Lévy-driven stochastic process under the risk-neutral measure induced by the Esscher transform. In the following section we work on a complete probability space $(\Omega, \mathcal{F}, P)$.

## 6.1. Pricing methodology

Classical asset pricing theory is based on the assumption of a complete market, where the risk associated with any derivative can be completely hedged against by replicating the derivative through a portfolio that includes holdings of the underlying asset. Then the derivative has a unique fair price equal to that of the replicating portfolio, and we say that the market is complete. This price can also be specified as the expected final payoff of the derivative under an equivalent measure $Q$, called the risk-neutral measure. Under this measure the discounted price processes of all tradeable assets are martingales.

For rainfall derivatives the underlying "asset" is an index $(I(t))_{t\geq 0}$ measuring accumulated rainfall, which cannot be directly traded, and so the hedging argument cannot be applied. Thus the market for rainfall derivatives is incomplete, meaning that there is no unique fair price of the derivative. Hence there exist many equivalent measures such that the discounted price processes of tradeable assets are martingales. In the present paper we construct one such measure by using the Esscher transform on the underlying rainfall process, which we specify through our Lévy-driven rainfall model.

The Esscher transform is a generalised Girsanov transform for jump processes; it was first introduced by Esscher (1932) as a change of probability measure, and Gerber and Shiu (1994) generalised the transform to stochastic processes driven by a Lévy process. As shown in Esche and Schweizer (2005), the Esscher transform preserves the Lévy properties of the process to be transformed. This property makes it a natural choice for constructing a risk-neutral measure when the underlying is driven by a Lévy process.

In the following we consider a finite time horizon $T < \infty$, and assume all derivatives expire before that time. We will use the generalised version of the Esscher transform for a Lévy process $(L(t))_{t\geq 0}$ with filtration $(\mathcal{F}_t)_{t\geq 0}$, which is defined by giving the Radon-Nikodym derivative

$$\left.\frac{dQ}{dP}\right|_{\mathcal{F}_t} = Z(t) = \frac{\exp\left\{\int_0^t \theta(s)dL(s)\right\}}{E\left[\exp\{\int_0^t \theta(s)dL(s)\}\right]}, \tag{12}$$

where $\theta(s)$ is a time-dependent parameter, as opposed to the standard transform where it is constant. This parameter can be interpreted as a measure of risk-aversion, called the *market price of risk*, and is used to calibrate $Q$ such that theoretical and observed market prices match. Specifically, the investor selling a derivative at time $t$ will have to pay an amount given by the payoff function of the index $I$ at the time of maturity $\tau$. This amount is determined by the jumps of the driving process $L$ in the future time interval $[t, \tau]$. Thus the investor is exposed to risk from these jumps, and the Esscher transform reflects the corresponding risk premium by exponentially tilting the jump measure.

Having defined $Q$ via the Esscher transform, we find derivative prices by taking expected values of payoffs at maturity under $Q$, conditional on the information known at the current time, similar to the complete market case. For simplicity we assume a zero interest rate. Then for a rainfall index $I$ adapted to the filtration $(\mathcal{F}_t)_{t\geq 0}$, we get that the futures price

process given by

$$f_\tau(t) = E_Q[I(\tau)|\mathcal{F}_t],$$

will be a $Q$-martingale by construction, which is required since the derivative contract is itself a tradeable asset. However, since the market is incomplete we do not require the underlying rainfall index process $(I(t))_{t\geq 0}$ to be a $Q$-martingale, since it cannot be directly traded.

### 6.2. Esscher transform for integrated moving average processes

In this subsection we show the result of applying the Esscher transform to the class of moving average processes, which includes our rainfall model as a special case.

Consider a Lévy subordinator $(L(t))_{t\geq 0}$. We can extend $L$ to a two-sided Lévy process $(L^*(t))_{t\in\mathbb{R}}$ by defining the process $L^*$ to be an independent copy of $L$ such that $L$ and $L^*$ have the same characteristic triplet, and letting

$$L^*(t) = \begin{cases} L(t), & \text{for } t \geq 0 \\ -L(-(t-)), & \text{for } t < 0, \end{cases}$$

which makes $L^*$ càdlàg. In the following we will take $L$ to mean $L^*$ in order to simplify notation. Also, we denote by $(F_t)_{t\in\mathbb{R}}$ the filtration associated with $(L(t))_{t\in\mathbb{R}}$.

We now define the stochastic process $X(t)$ by

$$X(s) = \int_{-\infty}^{s} h(s-v)\, dL(v) = \underbrace{\int_{-\infty}^{0} h(s-v)\, dL(v)}_{:=\tilde{A}(s)} + \int_0^s h(s-v)\, dL(v),$$

where $h : \mathbb{R}_+ \to \mathbb{R}_+$ is a left-continuous, square integrable deterministic function. The resulting process $X$ is strictly stationary, and is called a *moving average process* (Applebaum, 2009). A moving average process can be seen as a general form of the Ornstein-Uhlenbeck process, for the OU equation given by (2) we have $h(s) = \exp\{-\lambda s\}$, with $\tilde{A}(s) = e^{-\lambda s}X(0)$.

If we now integrate $X$ over the interval $[0,t]$ and exchange the order of integration by using the stochastic Fubini theorem, we get the integrated moving average process, which is similar to the primary rainfall model given in (3):

$$\begin{aligned} Y(t) &= \int_0^t X(s)\, ds = \underbrace{\int_{-\infty}^{0} \tilde{g}(t,v)\, dL(v)}_{:=A(0,t)} + \int_0^t g(t,v)\, dL(v), \\ g(t,v) &= \int_v^t h(s-v)\, ds, \quad \tilde{g}(t,v) = \int_0^t h(s-v)\, ds, \end{aligned} \tag{13}$$

with $A(0,t) \in \mathcal{F}_0 \ \forall t$.

We now want to calculate the characteristic function of the process $Y$ under the probability measure $Q$ specified by the Esscher transform defined in (12). In order to ensure that the Radon-Nikodym derivative $Z$ is well-defined, we assume that $L$ satisfies the exponential moment condition, which states that there exists a constant $k > 0$ such that

$$E[\exp(kL(t))] < \infty, \tag{14}$$

for $t < T$, where $T$ is our time horizon. For the particular case given by our rainfall model, $L$ is the Hougaard Lévy process $L(\mu, \phi, p)$, which has exponential moments for $k < \mu^{1-p}/(\phi(p-1))$.

As discussed in Subsection 6.1, derivative prices are calculated in terms of the expected payoff at maturity under the measure $Q$, conditional on the current information $(\mathcal{F}_t)$. We want to find prices for a general payoff function $f(\text{Ind}(\tau_1, \tau_2))$, where $\text{Ind}(\tau_1, \tau_2) = Y(\tau_2) - Y(\tau_1)$ is the index measuring accumulated rainfall in the interval $[\tau_1, \tau_2]$. We follow Benth and Šaltytė Benth (2013) in using Fourier methods for these calculations, where we define the Fourier transform and its inverse by

$$\begin{aligned} \hat{f}(y) &= \int_{\mathbb{R}} f(x) e^{-ixy}\, dx, \\ f(x) &= \frac{1}{2\pi} \int_{\mathbb{R}} \hat{f}(y) e^{ixy}\, dy, \end{aligned} \tag{15}$$

assuming $f, \hat{f} \in L^1(\mathbb{R})$. To account for the case when $f$ is not integrable, we consider the function $f_\delta(x) = e^{-\delta x} f(x)$, and assume there exists a $\delta \geq 0$ such that $f_\delta$ is integrable. We then have the following result specifying the risk-neutral expectation of $f(Y(\tau_2) - Y(\tau_1))$.

**Proposition 2.** *Let $f$ be a payoff function such that $\hat{f}_\delta \in L^1$, where $\hat{f}$ is the Fourier transform of $f$ as defined in* (15)*. Assume that*

$$\sup_{v \in [t, \tau_2]} (\delta |g(\tau_2, v) - g(\tau_1, v)| + |\theta(v)|) < k,$$

*for $k$ given by the exponential moment condition. Then, when $Y$ is an integrated moving average process as specified in* (13)*, and $Q$ corresponds to the Esscher transform, we have that*

$$\begin{aligned} &E_Q\left(f\left(Y(\tau_2) - Y(\tau_1)\right) | \mathcal{F}_t\right) \\ &= \frac{1}{2\pi} \int_{\mathbb{R}} \hat{f}_\delta(\xi) \exp\left\{ (\delta + i\xi) \left( A(0, \tau_2) - A(0, \tau_1) + \int_0^t [g(\tau_2, v) - g(\tau_1, v)]\, dL(v) \right) \right. \\ &\left. + \int_t^{\tau_1} \psi_\theta(v, (\delta + i\xi)[g(\tau_2, v) - g(\tau_1, v)])\, dv + \int_{\tau_1}^{\tau_2} \psi_\theta(v, (\delta + i\xi) g(\tau_2, v))\, dv \right\} d\xi, \end{aligned} \tag{16}$$

*for $t < \tau_1 < \tau_2 < T$, where we define*

$$\psi_\theta(s, \gamma c(s)) := \int_{\mathbb{R}_+} e^{\theta(s)y} (e^{(i\gamma c(s)y} - 1) \nu(dy),$$

*for a complex variable $\gamma$ and real-valued function $c$, where $\nu(\cdot)$ is the Lévy measure of $L$.*

The proof of Proposition 2 is given in the Appendix.

We stress that the resulting derivative price does depend on $t, \tau_1, \tau_2$, and as such our model constitutes an important generalisation of the independent increment model considered in Benth and Šaltytė Benth (2013). In their setup, the resulting futures prices did not exhibit any dynamics in $t$, which is a serious restriction in practical applications.

### 6.3. Pricing illustration

We now illustrate how to calculate the prices for futures written on the Detroit rainfall by using the model presented in Section 3, with the estimated parameters given in Table 3. For rainfall futures the payoff equals the index $Y(\tau_2) - Y(\tau_1)$, hence such contracts are also called swap contracts.

Recall that for the Detroit rainfall we used an integrated CARMA model of order $p = 1$, which reduces to the integrated OU model. This model can be written in the general form given in (13) above, with

$$g(t,s) = \frac{1 - e^{-\lambda(t-s)}}{\lambda}.$$

Letting $f(x) = e^{iux}$, we obtain the characteristic function $E_Q(e^{iu(Y(\tau-2)-Y(\tau_1))}|\mathcal{F}_t)$ as the exponential term in (16) evaluated at $\delta = 0$ and $\xi = u$. Taking derivatives with respect to $u$ and evaluating at $u = 0$ gives

$$E_Q\left[Y(\tau_2) - Y(\tau_1) \,|\, \mathcal{F}_t\right] = A(0,\tau_2) - A(0,\tau_1) + \left(\frac{e^{-\lambda\tau_1} - e^{-\lambda\tau_2}}{\lambda}\right)\int_0^t e^{\lambda v}\, dL(v)$$
$$+ \int_t^{\tau_1} -i\psi'_\theta(v,0)\left(\frac{e^{-\lambda\tau_1} - e^{-\lambda\tau_2}}{\lambda}\right) dv + \int_{\tau_1}^{\tau_2} -i\psi'_\theta(v,0)\left(\frac{1 - e^{-\lambda(\tau_2 - v)}}{\lambda}\right) dv,$$

where $\psi'_\theta$ denotes the derivative with respect to the second argument. The explicit form of $\psi_\theta$ for this model can be found analytically in terms of the parameters of the Hougaard process, and is specified in the Appendix. When $L$ is a Hougaard process the exponential moment condition (14) restricts the values of $\theta(v)$ to be below $\mu^{1-p}/(\phi(p-1)$, which equals 0.11 when using the estimated parameters. Note that this restriction does not affect the range of the prices, as the price explodes when $\theta(v)$ approaches the upper limit.

For our rainfall data we need to evaluate this expression based on the discrete observations $(Y(t_i))$, meaning that we do not observe $L$ or $A(0,\tau_i)$ directly. To obtain an explicit value for the price we approximate these terms by their expected value. It can be shown that for larger values of $\tau_1 - t$ the unobserved terms are negligible compared to the last two terms; hence the mean approximation does not significantly affect the value of the price.

We calculated prices for monthly rainfall contracts in 2011 for a time $t$ corresponding to the 31st of December 2010. The final price equals

$$S_{m_i}(t) E_Q[Y(\tau_2) - Y(\tau_1) \,|\, \mathcal{F}_t],$$

where $S_{m_i}(t)$ is the monthly average of the seasonality function $S(t)$. Table 6 shows a range of the resulting prices corresponding to different values of the parameter $\theta$, representing the risk premium.

The first row of the table shows market prices reported by the CME, and we see that by adjusting $\theta$ we can calibrate the prices obtained from the model to match the market price. The values of $\theta$ resulting from this calibration are shown in Table 7. We remark that rainfall derivatives are a very recent addition to the CME portfolio, and their trading volume is currently close to zero, thus the reported CME prices for 2011 do not accurately reflect the market value of these products at the current time. Hence the corresponding estimates of $\theta$ for the specific 2011 prices may also differ from their true value. These values do however demonstrate how the rainfall model and associated pricing methodology provides a unified and flexible framework for studying the market view of the risk associated with rainfall.

# 7. Conclusion

We have introduced a new class of continuous-time stochastic processes, driven by the Hougaard Lévy process, and shown how it can be used to construct a parsimonious and analytically tractable model for rainfall. By generalising the Ornstein-Uhlenbeck process

| | $\theta$ | Mar 11 | Apr 11 | May 11 | Jun 11 | Jul 11 | Aug 11 | Sep 11 | Oct 11 |
|---|---|---|---|---|---|---|---|---|---|
| CME price | | 4.20 | 4.40 | 3.20 | 5.0 | 4.5 | 4.3 | 4.2 | 4.6 |
| Expected price | -0.025 | 1.26 | 1.47 | 1.70 | 1.69 | 1.72 | 1.70 | 1.66 | 1.69 |
| | 0.000 | 1.76 | 2.06 | 2.39 | 2.37 | 2.40 | 2.38 | 2.33 | 2.36 |
| | 0.025 | 2.69 | 3.15 | 3.64 | 3.61 | 3.66 | 3.63 | 3.55 | 3.60 |
| | 0.050 | 4.74 | 5.54 | 6.42 | 6.37 | 6.46 | 6.41 | 6.26 | 6.35 |
| | 0.075 | 11.27 | 13.18 | 15.26 | 15.14 | 15.35 | 15.23 | 14.90 | 15.09 |
| | 0.100 | 76.13 | 89.08 | 103.07 | 102.27 | 103.72 | 102.92 | 100.65 | 101.95 |

Table 6: Prices of monthly rainfall contracts for Detroit.

| Month | Mar 11 | Apr 11 | May 11 | Jun 11 | Jul 11 | Aug 11 | Sep 11 | Oct 11 |
|---|---|---|---|---|---|---|---|---|
| $\theta$ | 0.045 | 0.041 | 0.018 | 0.040 | 0.035 | 0.033 | 0.033 | 0.037 |

Table 7: Estimated values of parameter $\theta$ based on CME prices.

representing rainfall intensity to a continuous-time ARMA (CARMA) process, we obtain a model with a very flexible autocorrelation structure. We presented a general fitting method for this class which exploits a correspondence between integrated CARMA and ARMA processes.

We showed that the model fits the marginal distribution of the rainfall very well on both hourly and daily time scales. In particular, the marginal fit for daily rainfall is better than that of the standard model described in Wilks (1998). The path properties of simulations from the model are similar to those of the empirical data, especially for data on a daily time scale.

By virtue of the CARMA generalisation, the extended model manages to accurately reproduce the autocorrelation structure of the observed rainfall, a characterising feature of the process which becomes increasingly significant for smaller time scales. We also showed how the ARMA correspondence can be used to obtain short-term predictions of the hourly rainfall.

The last part of the paper gives a result specifying the risk-neutral expectation of a function of the rainfall process, which can be used for pricing general derivatives written on a precipitation index. To construct a risk-neutral measure we use the Esscher transform, with a time-dependent parameter representing the risk premium. We state the result for a general moving average process, a class which includes our model as a special case. We illustrated our result by calculating futures prices based on empirical daily rainfall data from Detroit, and showed how they can be calibrated to observed prices. The pricing methodology constitutes an important generalisation of the independent increment model considered in Benth and Šaltytė Benth (2013), which does not allow for price dynamics of derivatives.

**Acknowledgements.** R. C. Noven gratefully acknowledges financial support from the Grantham Institute for Climate Change, Imperial College London. We thank the UK Meteorological Office and the British Atmospheric Data Centre for providing the data used.

# Appendix

In the following we present the proofs of our theoretical results.

First we quote a result (Cont and Tankov, 2004, Lemma 15.1) which will be used repeatedly in the proofs to follow:

**Lemma 1.** *Let $f : [0, T] \to \mathbb{R}$ be a left-continuous function and $L(t)$ a Lévy process. Then*

$$E\left[\exp\left\{\int_0^t if(s)dL(s)\right\}\right] = \exp\left\{\int_0^t \psi(f(s))ds,\right\}$$

*where $\psi(t)$ is the characteristic exponent of $L$, given by*

$$\exp\{\psi(u)\} = E[e^{iuL(1)}].$$

**Proof of Proposition 1.** The calculation of the covariance of $\triangle Y$ is based on the representation

$$\begin{aligned}\triangle Y(t) = &\int_0^\infty \left(\frac{e^{-\lambda t_{i-1}} - e^{-\lambda t_i}}{\lambda}\right) e^{-\lambda s}\, dL^*(s) + \int_0^{t_{i-1}} \left(\frac{e^{-\lambda(t_{i-1}-s)} - e^{-\lambda(t_i-s)}}{\lambda}\right) dL(s) \\ &+ \int_{t_{i-1}}^{t_i} \left(\frac{1 - e^{-\lambda(t_i-s)}}{\lambda}\right) dL(s),\end{aligned} \tag{17}$$

where $L^*$ is independent of $L$, and $L, L^*$ have the same characteristic triplet given by $(0, 0, \nu(\cdot))$, corresponding to the Hougaard process as specified in (4).

We first consider the following isometry (Klebaner, 2012, p.218):

$$E\left(\int_0^t H(s)\, dM(s)\right)^2 = E\left(\int_0^t H^2(s)\, d[M, M](s)\right), \tag{18}$$

for $(H(t))_{t\geq 0}$ a predictable process and $(M(t))_{t\geq 0}$ a square-integrable martingale. Now define the martingale $(M(t))_{t\geq 0}$ by

$$M(t) = L(t) - \mu t.$$

As $L$ is a pure-jump process we have

$$\begin{aligned}E([M, M](t)) = E([L, L](t)) &= E\left(\sum_{0\leq s\leq t} (\triangle L(s))^2\right) \\ &= t\int_0^\infty y^2\nu(dy) = t\text{Var[L(1)]} = t\phi\mu^p,\end{aligned} \tag{19}$$

where we used the Lévy-Khintchine formula for the subordinator case, which takes the form

$$E(e^{iuL(t)}) = \exp\{t\psi(u)\} = \exp\left\{t\int_{\mathbb{R}_+} (e^{iuy} - 1)\nu(dy)\right\}, \tag{20}$$

where $\nu(\cdot)$ is the Lévy measure associated with $L$. Combining (18) and (19), we have that for a deterministic, square-integrable function $H$,

$$\text{Var}\left(\int_0^t H(s)\, dL(s)\right) = E\left(\int_0^t H(s)\, dM(s)\right)^2 = \phi\mu^p \int_0^t H^2(s)\, ds.$$

We can now apply the above equality to find the covariance of the process $(\triangle Y(t))_{t\geq 0}$ as specified in (17), which gives the required result.

**Characteristic function of $\triangle Y$.** To obtain the characteristic function of $\triangle Y$ we first write it in the form given in (17), where the three terms are independent due to the independent

increments of $L$. We can then apply Lemma 1 to each term in the expression for the characteristic function. This gives an expression involving three integrals with integrands of form $\psi(f(s))$, which can be evaluated numerically in terms of the parameters of $L$.

**Proof of Proposition 2.** By construction of $f_\delta$, we have that

$$f(x) = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}_\delta(\xi) e^{(\delta+i\xi)x}\, d\xi,$$

and hence by the Fubini theorem, it follows that

$$E_Q[f(Y(\tau_2)-Y(\tau_1))|\,\mathcal{F}_t] = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}_\delta(y) E_Q[e^{(\delta+i\xi)(Y(\tau_2)-Y(\tau_1))}\Big|\,\mathcal{F}_t]\, d\xi,$$

similar to the proof of Proposition 8.4 in Benth and Šaltytė Benth (2013).

We now calculate the expectation involving the integrated moving average process $Y$. To this end, we first split the integrals in the expression for $Y$ as follows:

$$\begin{aligned}
&E_Q\left[\exp\left\{(\delta+i\xi)(Y(\tau_2)-Y(\tau_1))\right\}|\,\mathcal{F}_t\right] \\
=&\exp\left\{(\delta+i\xi)\left(A(0,\tau_2)-A(0,\tau_1)+\int_0^t [g(\tau_2,v)-g(\tau_1,v)]\, dL(v)\right)\right\} \\
\times&\underbrace{E_Q\left[\exp\left\{(\delta+i\xi)\int_t^{\tau_1}[g(\tau_2,v)-g(\tau_1,v)]\, dL(v)\right\}\Big|\,\mathcal{F}_t\right]}_{(\mathbf{A})} \\
&\times\underbrace{E_Q\left[\exp\left\{(\delta+i\xi)\int_{\tau_1}^{\tau_2} g(\tau_2,v)\, dL(v)\right\}\Big|\,\mathcal{F}_t\right]}_{(\mathbf{B})}.
\end{aligned} \tag{21}$$

By the abstract Bayes formula (see e.g Øksendal, 2000), for the measure $Q$ such that $dQ/dP|_{\mathcal{F}_t} = Z(t)$, with $X$ being $\mathcal{F}_\tau$-measurable and $t<\tau$, we have that

$$E_Q(X|\mathcal{F}_t) = E\left(X\frac{Z(\tau)}{Z(t)}\Big|\,\mathcal{F}_t\right).$$

Recall that we are working with the Esscher transform, so we have

$$Z(t) = \frac{\exp\left\{\int_0^t \theta(v)\, dL(v)\right\}}{E\left[\exp\left\{\int_0^t \theta(v)\, dL(v)\right\}\right]}.$$

Applying the Esscher transform then gives

$$\begin{aligned}
(\mathbf{A}) =& E\left[\exp\left\{\int_t^{\tau_1}(\delta+i\xi)\left[g(\tau_2,v)-g(\tau_1,v)\right]\, dL(v)\right\}\frac{Z(\tau_1)}{Z(t)}\Big|\,\mathcal{F}_t\right] \\
=& E\left[\exp\left\{\int_t^{\tau_1}\left((\delta+i\xi)\left[g(\tau_2,v)-g(\tau_1,v)\right]+\theta(v)\right)\, dL(v)\right\}\right]\times\frac{E\left[\exp\left\{\int_0^t\theta(v)\, dL(v)\right\}\right]}{E\left[\exp\left\{\int_0^{\tau_1}\theta(v)\, dL(v)\right\}\right]},
\end{aligned}$$

where we get an unconditional expectation due to the independent increments of $L$.

We can extend Lemma 1 to complex-valued functions to get

$$E\left[\exp\left\{\int_0^t (a(v)+ib(v))\, dL(v)\right\}\right] = \exp\left\{\int_0^t \psi_1(-ia(v)+b(v))\, dv\right\}, \tag{22}$$

where the term on the RHS equals

$$\exp\left\{\int_0^t \ln E\left(\exp\{a(v)+ib(v)\}L(1)\right)\,dv\right\},$$

and we have that

$$\left|E\left(e^{(a(v)+ib(v))L(1)}\right)\right| \leq E(e^{a(v)L(1)}),$$

and so if $\sup_v a(v) < k$, then the last term is bounded by the exponential moment condition given in (14).

Applying (22) to the terms in (**A**) gives

$$(\mathbf{A}) = \exp\left\{\int_t^{\tau_1} \psi\Big(\xi\left[g(\tau_2,v)-g(\tau_1,v)\right] - i\left[\theta(v)+\delta\left(g(\tau_2,v)-g(\tau_1,v)\right)\right]\Big) - \psi\left(-i\theta(v)\right)\,dv\right\},$$

where the requirement $\sup_{v\in[t,\tau_2]}\left(|\theta(v)|+\delta|g(\tau_2,v)-g(\tau_1,v)|\right) < k$ ensures that the terms in the above equation are well-defined.

Now we consider the Lévy-Khintchine formula for subordinators, given in (20), where $\nu$ is the Lévy measure of $L$. We can analytically continue this formula to complex arguments (Applebaum, 2009, p. 338), and so we get that

$$(\mathbf{A}) = \exp\left\{\int_t^{\tau_1}\int_{\mathbb{R}_+} e^{\theta(v)y}(e^{(\delta+i\xi)[g(\tau_2,v)-g(\tau_1,v)]y}-1)\nu(dy)\,dv\right\}.$$

We note that this expression can also be written as

$$E\left[\exp\left\{\int_t^{\tau_1}(\delta+i\xi)\left[g(\tau_2,v)-g(\tau_1,v)\right]\,dL_Q(v)\right\}\right],$$

where $L_Q(v)$ is now a non-stationary stochastic process with jump measure depending on time, namely

$$\nu_\theta(dv,dy) = e^{\theta(v)y}\nu(dy)dv.$$

Thus we see that conditioning with respect to the measure $Q$ has the effect of exponentially tilting the jump measure of $L$ at time $v$ according to $\theta(v)$, so the jumps of $L$ at times $v$ will be weighted more or less in the expectation depending on the sign of $\theta(v)$.

Now defining

$$\psi_\theta(v,\gamma c(\cdot)) := \int_{\mathbb{R}_+} e^{\theta(v)y}(e^{\gamma c(v)y}-1)\nu(dy), \tag{23}$$

we get that

$$(\mathbf{A}) = \exp\left\{\int_t^{\tau_1}\psi_\theta\Big(v,(\delta+i\xi)\left[g(\tau_2,v)-g(\tau_1,v)\right]\Big)\,dv\right\};$$

and by similar arguments

$$(\mathbf{B}) = \exp\left\{\int_{\tau_1}^{\tau_2}\psi_\theta\Big(v,(\delta+i\xi)g(\tau_2,v)\Big)\,dv\right\}.$$

Substituting these expressions into (21) then gives the result.

**Hougaard process.** The Hougaard process has Lévy measure given by (Grigelionis, 2011)

$$\nu(dy) = \left(\phi^{\frac{1}{p-1}}\Gamma\left(\frac{p}{p-1}\right)(p-1)^{p/(p-1)}\right)^{-1} y^{\frac{3-2p}{p-1}} \exp\left\{-\frac{\mu^{1-p}}{\phi(p-1)}y\right\} dy,$$

in terms of the Tweedie parameterisation. We also have that when $L$ is the Hougaard process, the function $\psi_\theta$ defined in (23) takes the form

$$\psi_\theta(v, \gamma c(v)) = \frac{\mu^{2-p}}{\phi(2-p)}\left[\left(1 - \frac{\phi(p-1)\left(i\gamma c(v) + \theta(s)\right)}{\mu^{1-p}}\right)^{(p-2)/(p-1)} - 1\right].$$

# References


Applebaum, D. (2009), *Lévy Processes and Stochastic Calculus*, second edn, Cambridge University Press.

Barndorff-Nielsen, O. E. and Shephard, N. (2001), 'Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics', *Journal of the Royal Statistical Society - Series B: Statistical Methodology* **63**(2), 167–241.

Benth, F. and Šaltytė Benth, J. (2013), *Modeling and pricing in financial markets for weather derivatives*, first edn, World Scientific Publishing.

Brockwell, P. (2001), 'Lévy-Driven CARMA Processes', *Annals of the Institute of Statistical Mathematics.* **53**(1), 113–124.

Brockwell, P. (2004), 'Representations of continuous-time ARMA processes', *Journal of Applied Probability* **41**(2004), 375–382.

Brockwell, P. J. and Lindner, A. (2009), 'Existence and uniqueness of stationary Lévy-driven CARMA processes', *Stochastic Processes and their Applications* **119**(8), 2660–2681.

Brockwell, P. and Lindner, A. (2013), 'Integration of CARMA processes and spot volatility modelling', *Journal of Time Series Analysis* **34**(2), 156–167.

Carmona, R. and Diko, P. (2005), 'Pricing Precipitation Based Derivatives', *International Journal of Theoretical and Applied Finance* **08**(07), 959–988.

Chin, E. H. (1977), 'Modeling daily precipitation occurrence process with Markov Chain', *Water Resources Research* **13**(6), 949–956.

Coe, R. and Stern, R. (1982), 'Fitting models to daily rainfall data', *Journal of Applied Meteorology* **21**, 1024–1031.

Cont, R. and Tankov, P. (2004), *Financial modelling with jump processes*, Chapman & Hall/CRC.

Cowpertwait, P., Isham, V. and Onof, C. (2007), 'Point process models of rainfall: developments for fine-scale structure', *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **463**(2086), 2569–2587.

Cox, D. R. and Isham, V. (1988), 'A Simple Spatial-Temporal Model of Rainfall', *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **415**(1849), 317–328.

Diebold, F. and Mariano, R. (1995), 'Comparing predictive accuracy', *Journal of Business & Economic Statistics* **13**(3), 253–263.

Dunn, P. K. (2004), 'Occurrence and quantity of precipitation can be modelled simultaneously', *International Journal of Climatology* **24**(10), 1231–1239.

Esche, F. and Schweizer, M. (2005), 'Minimal entropy preserves the Lévy property: how and why', *Stochastic Processes and their Applications* **115**(2), 299–327.

Esscher, F. (1932), 'On the probability function in the collective theory of risk', *Scandinavian Actuarial Journal* **15**.

Gerber, H. and Shiu, E. (1994), 'Option pricing by Esscher transforms', *Transactions of the Society of Actuaries* **46**.

Grigelionis, B. (2011), 'On the Hougaard subordinated Gaussian Lévy processes', *Statistics & Probability Letters* **81**(8), 998–1002.

Jørgensen, B. (1997), *The Theory of Dispersion models*, Chapmann&Hall.

Katz, R. (1977), 'Precipitation as a chain-dependent process', *Journal of Applied Meteorology* .

Klebaner, F. (2012), *Introduction to Stochastic Calculus with Applications*, third edn, Imperial College Press.

Kunsch, H. R. (1989), 'The Jackknife and the Bootstrap for General Stationary Observations', *The Annals of Statistics* **17**(3), 1217–1241.

Lee, M. and Whitmore, G. (1993), 'Stochastic processes directed by randomized time', *Journal of applied probability* **30**(2), 302–314.

Leobacher, G. and Ngare, P. (2011), 'On modelling and pricing rainfall derivatives with seasonality', *Applied Mathematical Finance* **18**(1), 1–24.

López Cabrera, B., Odening, M. and Ritter, M. (2013), 'Pricing rainfall futures at the CME', *Journal of Banking & Finance* **37**(11), 4286–4298.

Northrop, P. (1998), 'A clustered spatial-temporal model of rainfall', *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **454**(1975), 1875–1888.

Øksendal, B. (2000), *Stochastic Differential Equations: an introduction with applications*, second edn, Springer-Verlag Berlin.

Onof, C., Chandler, R. E., Kakou, A., Northrop, P., Wheater, H. S. and Isham, V. (2000), 'Rainfall modelling using Poisson-cluster processes: a review of developments', *Stochastic Environmental Research and Risk Assessment* **14**(6), 384 – 411.

Politis, D. N. and Romano, J. P. (1994), 'The Stationary Bootstrap', *Journal of the American Statistical Association* **89**(428), pp. 1303–1313.

Rodriguez-Iturbe, I., Cox, D. R. and Isham, V. (1987), 'Some models for rainfall based on stochastic point processes', *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **410**(1839), 269–288.

UK Meteorological Office (2012), 'Met Office Integrated Data Archive System (MIDAS) Land and Marine Surface Stations Data (1853-current)'. NCAS British Atmospheric Data Centre. Available from `http://badc.nerc.ac.uk/view/badc.nerc.ac.uk__ATOM__dataent_ukmo-midas`.

Wilks, D. (1998), 'Multisite generalization of a daily stochastic precipitation generation model', *Journal of Hydrology* **210**(1-4), 178–191.

Woolhiser, D. A. and Roldán, J. (1982), 'Stochastic daily precipitation models: 2. A comparison of distributions of amounts', *Water Resources Research* **18**(5), 1461–1468.